

REMARKS

Claims 29-33, 38, 43-45, and 49-51 are pending in the application.

Claims 29-33, 38, 43-45, and 49-51 are currently amended and claims 36, 39, 47, and 56 are canceled. Applicants respectfully submit that no new matter is added to currently amended claims 29-33, 38, 43-45, and 49-51.

Claims 29-33, 36, 38-39, 43-45, 47, 49-51, and 56 stand rejected under 35 U.S.C. §103(a) as unpatentable over U.S. Patent Application Publication No. 2004/0015386 to Abe et al., hereinafter, Abe in view of U.S. Patent Application Publication No. 2004/0117239 to Mittal et al., hereinafter, Mittal.

Applicant respectfully traverses the rejections based on the following discussion. The following paragraphs are numbered for ease of future reference.

I. The 35 U.S.C. 103(a) Rejection over Abe and Mittal

A. The Abe Disclosure

[0001] It is a fact that Abe discloses, "A system and method for sequential decision-making for customer relationship management includes providing customer data including stimulus-response history data, and automatically generating actionable rules based on the customer data. Further, automatically generating actionable rules may include estimating a value function using reinforcement learning." (Abstract).

[0002] It is a fact that Abe discloses, "The present invention includes an inventive method for sequential decision making (e.g., sequential cost-sensitive decision making) for customer relationship management. The inventive method includes providing customer data (e.g., consumer data, client data, donor data, etc.) comprising stimulus-response history data, and automatically generating actionable rules based on the customer data. Further, automatically generating actionable rules may include estimating a value function using reinforcement learning (e.g., reinforcement learning and dynamic programming). For example, estimating a value function may include value iteration." (paragraph [0014]).

[0003] It is a fact that Abe discloses, "... the present invention may utilize the popular

Markov Decision Process model with function approximation. In a Markov Decision Process (MDP), the environment is assumed to be in some state at any given point in time. In the case of targeted marketing, such states would be represented as feature vectors comprising categorical and numerical data fields that characterize what is known about each customer at the time a decision is made." (paragraph [0066]).

[0004] It is a fact that Abe discloses, "When the learner takes an action, it receives a finite reward and the environment makes a probabilistic transition to another state. The goal of a learner is to learn to act so as to maximize the cumulative reward it receives (usually with future rewards discounted) as the learner takes actions and traverses through the state space. In the example of targeted marketing, a customer, with all her past history of purchases and promotions, is in a certain state at any given point in time. When a retailer takes an action, the customer then makes a probabilistic transition to another state, possibly generating a reward. ... The reward at each state transition is the net profit to the retailer. ... Application of reinforcement learning to this problem amounts to maximizing the net present value of profits and losses over the life cycle of a customer." (paragraph [0067]).

[0005] It is a fact that Abe discloses, "At any point in time, the environment is assumed to be in one of a set of possible states. At each time tick (the present invention may assume a discrete time clock), the environment is in some state s , the learner takes one of several possible actions a , receives a finite reward (i.e., a profit or loss) r , and the environment makes a transition to another state s' . Here, the reward r and the transition state s' are both obtained with probability distributions that depend on the state s and action a ." (paragraph [0074]).

[0006] It is a fact that Abe discloses, "The inventors also conducted experiments to examine the effect of using the various sampling methods proposed hereinabove with respect to the quality of the output models and the required computational resources. FIG. 12 plots the total life-time profits attained using different sampling methods as a function of the number of value iterations that were performed. The sampling methods employed were random sampling, Q-sampling, TD(λ)-sampling with 2-step lookahead, and TD(λ)-sampling with 3-step lookahead." (paragraph [0150]).

[0007] It is a fact that FIG. 12 of Abe discloses,




FIG.12 ITERATIONS

[0008] It is a fact that Abe discloses, "... the present invention includes a method for optimized sequential targeted marketing. The method may include preparing data, estimating a value function, and transforming rules. Data preparing may include using customer data such as demographic features, transaction history data such as purchase records, web, wireless, kiosk access data, call center records, which may be used to generate a sequence of event data, where each event datum consists of demographic features of a customer, if any, and a number of features of the same customer which collectively reflect the state of that customer at a certain point in time. Such features may be derived from the customer's transaction history data (e.g., the number of purchases made to date, number of purchases made in recent months, the amount of purchases made, the amount of purchases made recently, the frequency of web access, frequency of web access made recently, possibly categorized by the types of web pages, etc.). Such data may also include the marketing action which may be taken at or around that time by the seller (e.g., retailer), the response may be taken by that customer at or around that time, and the amount of profit or cost associated with that action, if available." (paragraph [0152]).

B. The Mittal Disclosure

[0009] It is a fact that Mittal discloses, "The invention describes a method and system for conducting online marketing research keeping in consideration the specified budget for the experiment. The invention describes a methodology for effective data collection and optimised utilisation of budget through the use of efficient sampling and grouping of users." (Abstract).

[0010] It is a fact that Mittal discloses, "The merchant desires to complete the experiment within the specified budget, within the defined time period and obtain the required information from one or more set of target users/customers." (paragraph [0015]).

C. Argument

[0011] It is a fact that the present invention discloses, "FIG. 4 is a flowchart depicting the reinforcement learning algorithm, as it exists in the art. Step 402 estimates an initial value, $Q'(s,a)$ for all states s and actions a . At step 404 an action a^* having the maximum estimated value of $Q'(s,a)$ for a given state s is identified. That is, $Q'(s,a') = \max_{a} Q'(s,a)$. At step 406, an action a' having deployment probability ϵ is chosen. Step 406 uses the following randomization to select an action in state s for deployment to enable access by customers." (Pub. No. 2005/0071223, paragraph [0084]).

[0012] It is a fact that FIG. 4 of the present invention discloses,

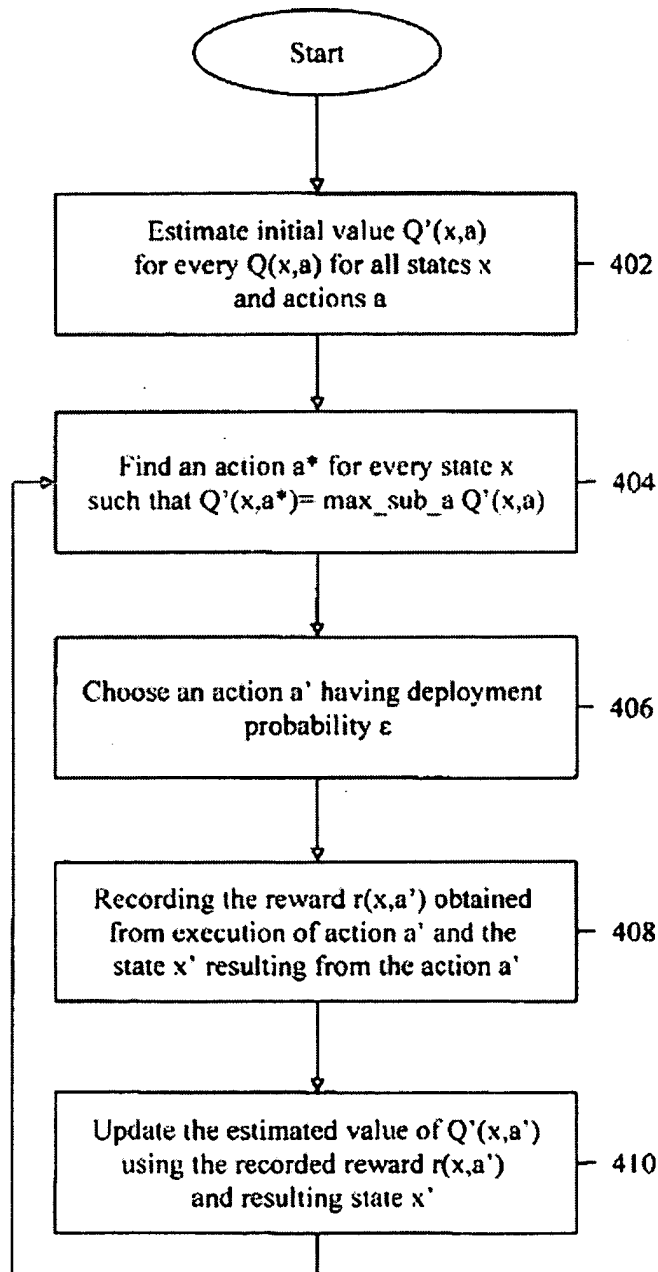


FIG. 4

[0013] It is a fact that the present invention discloses, "To allow for exploration of other actions, an action different from a^* suggested in the algorithm is selected occasionally. This is

done through some randomization. To draw an analogy, this randomization procedure can be viewed as tossing a biased coin (where heads and tails are not equally probable, rather head occurs with probability $1-\epsilon$ and tails with probability ϵ for some positive $\epsilon > 0$). The coin is unbiased if $\epsilon = 1/2$. If tail results in head, a^* is used in the execution. But if a toss results in tails, then any action (chosen arbitrarily or uniformly) other than a^* is used for execution." (Pub. No. 2005/0071223, paragraph [0085]).

[0014] It is a fact that the present invention discloses, "Corresponding to action a^* with probability $1-\epsilon$ another action a' with probability ϵ for some positive $\epsilon > 0$ is selected. The action, a' resulting from such randomization is then executed. At step 408 the reward $r(s, a')$ obtained from the execution of randomized action a' and the new state, s' , resulting from this action is recorded." (Pub. No. 2005/0071223, paragraph [0085]).

[0015] As supported by the Specification of the present invention above, the independent claims of the present invention clearly describe at least the features of: "recommending, by said computer, a set of possible marketing strategies along with a deployment probability of each of said set of possible marketing strategies to determine an optimal marketing strategy by using a modified Reinforcement Learning (RL) algorithm", as recited in currently amended, independent claims 29 and 49, and as similarly recited in currently amended, independent claim 43.

[0016] Nowhere does Abe disclose, teach or suggest a randomization of actions to allow for exploration of other actions in a RL algorithm as does the present invention. (Pub. No. 2005/0071223, paragraph [0085]).

[0017] Instead, Abe merely discloses a method of sequential decision-making for customer relationship management that includes providing customer data including stimulus-response history data, and generating actionable rules based on the customer data, where the actionable rules may include estimating a value function using reinforcement learning." (Abstract).

[0018] Therefore, Abe does not disclose, teach or suggest at least the present invention's features of: "recommending, by said computer, a set of possible marketing strategies along with a deployment probability of each of said set of possible marketing strategies to determine an optimal marketing strategy by using a modified Reinforcement Learning (RL) algorithm", as

recited in currently amended, independent claims 29 and 49, and as similarly recited in currently amended, independent claim 43.

[0019] Mittal does not cure the deficiencies of Abe.

[0020] Mittal merely discloses a method for effective data collection in online marketing research and optimised utilisation of a budget through the use of efficient sampling and grouping of users. (Abstract).

[0021] Nowhere does Mittal disclose, teach or suggest a randomization of actions to allow for exploration of other actions in a RL algorithm as does the present invention. (Pub. No. 2005/0071223, paragraph [0085]).

[0022] Therefore, Mittal does not disclose, teach or suggest at least the present invention's features of: "recommending, by said computer, a set of possible marketing strategies along with a deployment probability of each of said set of possible marketing strategies to determine an optimal marketing strategy by using a modified Reinforcement Learning (RL) algorithm", as recited in currently amended, independent claims 29 and 49, and as similarly recited in currently amended, independent claim 43.

[0023] Instead, Mittal merely discloses a method for effective data collection in online marketing research and optimised utilisation of a budget through the use of efficient sampling and grouping of users. (Abstract).

[0024] For at least the reasons outlined above, Applicants respectfully submit that Abe and Mittal, either individually or in combination, do not disclose, teach or suggest at least the present invention's features of: "recommending, by said computer, a set of possible marketing strategies along with a deployment probability of each of said set of possible marketing strategies to determine an optimal marketing strategy by using a modified Reinforcement Learning (RL) algorithm", as recited in currently amended, independent claims 29 and 49, and as similarly recited in currently amended, independent claim 43. Accordingly, Abe and Mittal, either individually or in combination, fail to render obvious the subject matter of currently amended, independent claims 29, 43, and 49, and currently amended, dependent claims 30-33, 38, 44, 45, 50, and 51 under 35 U.S.C. §103(a). The rejection of canceled claims 36, 39, 47, and 56 is moot.

Withdrawal of the rejection of claims 29-33, 36, 38-39, 43-45, 47, 49-51, and 56 under 35 U.S.C. §103(a) as unpatentable over Abe and Mittal is respectfully solicited.

II. Formal Matters and Conclusion

Claims 29-33, 38, 43-45, and 49-51 are pending in the application.

With respect to the rejection of the claims over the cited prior art, Applicants respectfully argue that the present claims are distinguishable over the prior art of record. In view of the foregoing, the Examiner is respectfully requested to reconsider and withdraw the rejections to the claims.

In view of the foregoing, Applicants submit that claims 29-33, 38, 43-45, and 49-51, all the claims presently pending in the application, are in condition for allowance. The Examiner is respectfully requested to pass the above application to issue at the earliest time possible.

Should the Examiner find the application to be other than in condition for allowance, the Examiner is requested to contact the undersigned at the local telephone number listed below to discuss any other changes deemed necessary.

Please charge any deficiencies and credit any overpayments to Attorney's Deposit Account Number 09-0441.

Respectfully submitted,

Dated: June 22, 2009

/Peter A. Balnave/
Peter A. Balnave, Ph.D.
Registration No. 46,199

Gibb I.P. Law Firm, LLC
2568-A Riva Road, Suite 304
Annapolis, MD 21401
Voice: (410) 573-5255
Fax: (301) 261-8825
Email: Balnave@gibbiplaw.com
Customer Number: 29154